The DARPA JFACC Program:
Modeling and Control of Military Operations

Sharon A. Heise[1] and H. Stephen Morse[2]
DARPA/JFACC
3701 N. Fairfax Drive
Arlington, VA 22203
email: saheise@darpa.mil, hmorse@snap.org

Abstract

In the Introduction, the goals, scope, and structure of the DARPA JFACC program will be described, with emphasis on the flexibility given to investigators in problem definition and research objectives, and the wide variety of approaches that have been proposed. In Section 2, five major areas of research will be briefly summarized: active adversary and game theoretic techniques; stochastic optimization; model predictive control; finite state machines and discrete event systems; and emergent behavior. Section 3 will describe the major Concepts of Operation under consideration for application of this research: fully autonomous operation; state estimation and prediction; tracking and adaptive control; and on-line sensitivity analysis. Section 4 will present significant conclusions arising out of current research, and future directions.

1. Introduction

In September of 1999, the DARPA JFACC program initiated work on a new approach to military command and control (C2). Up to that point, the program had been built around techniques from the Artificial Intelligence community – techniques typically using rule-based systems, tree structures and associated search algorithms, or pattern matching. This new approach is centered around the following hypothesis: *Control theory can make a significant contribution to selected problems in military command and control.* For the past year or so, several research teams from both academic institutions and contractor laboratories have been involved in testing and bounding this hypothesis. The purpose of this paper is to provide a brief introduction to some of their work.

Two concepts – *agility*, and *stability* – were used to focus and motivate the research. In the military context, *agility* is defined as rapid and optimal response within a highly dynamic battle environment. Similarly, *stability* is defined as the recognition that plan-to-plan similarity, or overlap, is itself of significant military value. A guiding metaphor has been the automatic flight control system of the F-14 fighter aircraft. This system permits routine and safe operation of the aircraft under flight conditions, which would be impossible for an unaided human being to cope with. By analogy, it is conjectured that control theory applied to the battlespace could provide a prosthesis, enabling commanders to respond with a level of timeliness and precision heretofore unachievable.

To test this hypothesis, each researcher performs the classic control experiment: a controller is paired against a plant – in this case, a software simulation of the battlespace. The diversity of modeling approaches and solution techniques proposed by the researchers did not permit independent development of a single plant simulation. Instead, each researcher has developed a matched plant, and a software audit has verified that double-bookkeeping of state variables in the plant and the controller has been maintained (see Figure 1).

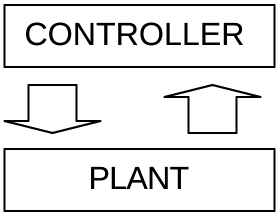

Figure 1.

As will be seen, the close coupling of the controller and its matched simulated plant has given rise to some interesting concepts of operation not originally envisioned (see Section 3 below).

The remainder of this paper is organized as follows. In Section 2, five areas of research within the project are briefly described. While these do not exhaust the approaches under investigation, they will give a good sense of the breadth of thinking that has been generated. Section 3 will turn to issues of

---

[1] Program Manager, Major, USAF

[2] Chief Technologist

application – that is, concepts of operation for the transition of these control technologies to military operations. A final section then highlights some key lessons learned, and indicates where the future may lead us.

2. Areas of Research

In this section five major areas of research among the JFACC program teams will be highlighted: game theoretic techniques with an active, intelligent adversary; stochastic optimization; model predictive control; discrete event systems; and emergent behavior.

While some researchers have built fairly elaborate event driven simulations similar to those used in military war games, most of the efforts described here have abstracted the battle down to a few key entities and objectives: a few types of aircraft (bombers, radar jammers, and attack fighters, for both friendly and opposition forces); opposition radar units, capable of detecting and firing on friendly aircraft; and a set of fixed or moving ground targets to be located, fired upon, and disabled. Values for assets and targets may be assigned and used in objective functions. There may also be models for detection (based on range and type of asset) and attrition (based on relative numbers and types of units involved in an engagement). A typical scenario is for friendly aircraft to penetrate a thicket of surface-to-air missile (SAM) sites, and to disable a set of fixed, known, and defended targets.

2.1 Active Intelligent Adversaries

Rather than treat the actions of the adversary as a disturbance, several researchers have modeled the problem as a game, in which the adversary has its own set of controls issued to its own controllable units in the jointly accessible plant. Here, the paradigm experiment shown in Figure 1 must be extended (see Figure 2):

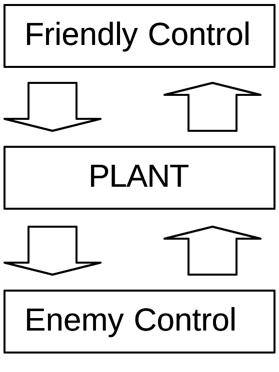

Figure 2.

Researchers began by assuming a zero-sum game (or an almost-zero-sum game) in which Blue and Red controllers share an objective function, each attempting to drive it in opposite directions. This gives rise to classic game solutions using the Nash equilibrium, or using Stackelberg leader-follower results. However, it was soon realized that this is often an unrealistic assumption. In fact, the goals of the adversary may be different from ones own, leading to unexpected adversary behaviors in the plant.

The next stage, then, was to treat the parameters of the objective function used by the adversary as *unknown state to be estimated*, based on observation of adversary action in the plant. This is a special form of the state estimation problem, and holds significant promise as a technique for inferring adversarial intent, and perhaps for exploiting that knowledge to one’s own advantage.

Another game theoretic technique under investigation treats the problem as a three-dimensional *abstract board game*, similar to chess. Moves of pieces are developed to mimic the dynamic constraints of the battlespace, and *Linguistic Geometry* [5] is used to rapidly construct highly optimized sequences of moves to achieve desired goals.

2.2 Stochastic Optimization

Several investigators are optimizing control strategies by considering explicit probabilistic models of uncertainty in the plant. The most common source of uncertainty concerns attrition of engagements, but other sources include uncertainty in detection of objects by sensors, and uncertainty in location and type of opposing assets.

One approach is to use detailed event driven simulations – often using multiple independent Monte Carlo runs – to estimate the conditional distributions and correlations of key parameters (for example, the likelihood of surviving an attack if a particular route is selected for attacking a target). These probabilities are then used in open-loop optimization algorithms employing, for example, linear or integer programming, or auction algorithms. A control strategy covering some time horizon (perhaps several hours) is selected and issued for execution. When the actual results of the attack become available, the state is updated and the process is repeated. We may consider this an example of repeated open loop control similar to the classic military OODA-Loop (observe, orient, decide, act).

A more sophisticated approach taken by some researchers is to solve the *closed loop* optimal stochastic control problem [1,4]. Dynamic programming is the algorithm of choice here, but (as is well known) this algorithm does not scale well to large problems without carefully selected heuristics and/or pruning techniques.

This approach also offers a possible solution to the problem of objective functions that rely on estimates for target value. Instead, one is able to consider objective functions based on likelihood, or expected time, to achieve a certain state, independent of target value. This is of operational interest since it is well known that estimation of target value is a difficult and often misleading surrogate for the commander's intent.

Several researchers have proposed strategies for dealing with non-linearities and coupling effects in targeting. Others have proposed objective functions that either do not rely on target value, or that only require knowledge of ordinal (rather than cardinal) relationships.

2.3 Model Predictive Control (MPC)

MPC techniques map naturally onto the JFACC problem domain and provide a powerful architectural approach capable of effectively utilizing and integrating a variety of modeling techniques. Returning to the original goals of agility and stability, MPC naturally addresses agility by recomputing the optimal solution (over its moving horizon) in response to newly arrived information. It can further address stability issues by including plan-to-plan variability (over the portion of the horizon shared by successive plans) as one term in its chosen objective function.

As is clear by comparing (for example) [1] and [4], there is a strong relationship between MPC and closed loop optimal stochastic control. This relationship is clearly shown in the various approaches taken by several researchers. The key idea in both approaches to make a control decision based on an optimization over an extended time horizon, and to then update that decision when and as new information becomes available. Viewed in this way, many different modeling techniques could contribute to an MPC-based solution, e.g. any modeling technique with the ability to extrapolate current options forward in time so as to compare likely outcomes.

Another research objective being considered within the MPC context is the linkage between optimizing control of weapons systems, on the one hand, and optimizing control of the placement of sensor resources, on the other hand. While the JFACC program has not explicitly included ISR (intelligence, surveillance, and reconnaissance) in its purview, there is clearly a strong linkage. A controller needs state feedback, and optimizing the collection of that feedback arises naturally in the MPC problem formulation. The artificial wall separating these issues in the current program will hopefully be torn down in future efforts.

2.4 Discrete Event Controllers

Some researchers have chosen to model entities, and their interactions, in terms of *discrete event systems* (or, equivalently, *finite state machines*, or *formal languages*). Starting with abstract states and allowed transitions, formal languages (that is, sequences of transitions) are defined that lead to desired goal states and/or avoid undesirable ones. When one DES sits above another, and issues "controls" that enable or disable selected state transitions, a Discrete Event Controller results [2,3].

Such a highly formal approach permits completeness proofs similar to a compiler for a computer language. A key open research issue is how the mapping between the upper and lower DES controllers should be realized.

DES's have been effectively used to implement Hybrid Controllers, in which an element from a collection of piece-wise linear control laws is selected by the DES based on threshold crossings of key global parameters. This marriage between the formality of DES models and the power of locally linear continuous models appears to be useful and widely applicable. The DES researchers in the JFACC program have included continuous models to deal with specific localized optimizations. This practice greatly reduces the size of the state space (a major hurdle for DES approaches), but increases the complexity of the level-to-level mapping.

DES control is of great interest to the JFACC program because of its potential applicability to robotics. If a future battlespace is imagined in which humans and robots synchronize activities within a shared geographic extent to achieve a common objective, a hierarchy of DES's appears to be a natural and powerful way to model the robotic component. A more ambitious research goal is to automatically generate level-to-level DES mappings based on the entropy of language sequences derived from actual combat data.

2.5 Emergent Behavior

As a final example of research, consider an approach that attempts to control swarms of simple entities in the battlespace using the biological metaphor of *insect pheromones*. The key idea here is that desired emergent behavior can be selectively generated by choosing such low-level and localized factors as types of pheromone, deposition rates, evaporation rates, and gradient following policies.

This approach is of interest for several reasons. As with the previous research area, emergent behavior has clear applicability to a battlespace filled with a large number of simple robotic entities. Just as important, however, is the ability to model highly non-linear effects, i.e. complexity. Even if the "swarms" are not instantiated as hardware, the behaviors they are able to generate via simulation may well-approximate important dynamics in the battlespace. Thus, emergent behavior could prove to be a critical modeling technique for prediction and sensitivity analysis.

The algorithmic and computational difficulties associated with system non-linearities are well known. Suppose that predictive modeling techniques employing emergent behaviors were coupled with an MPC-style control algorithm. What might the result look like? It is questions like this that are on the mind of the JFACC program researchers, and that make this project so interesting and volatile.

3. Concepts of Operation

In this section, four possible ways of utilizing the technologies that have been proposed by researchers will be examined: fully autonomous operations; state estimation and prediction; tracking and adaptive control; and sensitivity analysis.

3.1 Fully Autonomous Operation

In the traditional setting for control theory, the control system is deployed with no (or minimal) supervision, maintaining the safety of the system, and driving it toward some desired goal. In military operations scenarios, however, there are operational constraints that increase risk, and make acceptance of this mode of operation difficult. Perhaps the most likely application is in light, inexpensive sensors, where system failure would not lead to potentially disastrous results.

It should also be remembered that these controllers will exist in an embedded hierarchy, each level of which has its own appropriate level of abstraction and time scale. The further down this hierarchy one moves – from the strategic, to the operational, to the tactical – the greater the opportunities for largely unattended operation.

Both the DES approach and the emergent behavior approach are well-suited to this type of application. As the military moves toward larger numbers of unmanned or remotely controlled entities (largely as a way to reduce direct risk to the warfighter), these techniques will receive correspondingly increased attention.

3.2 State Estimation and Prediction

Each of the next three methods of application rest upon a particular way of utilizing the controller – not directly against the real plant (the battlespace), but against the simulated plant as illustrated in Figures 1 and 2. If this arrangement is thought of as a paired and integrated collection of software components separate from the real world, and linked to it by state updates, the situation shown in Figure 3 arises.

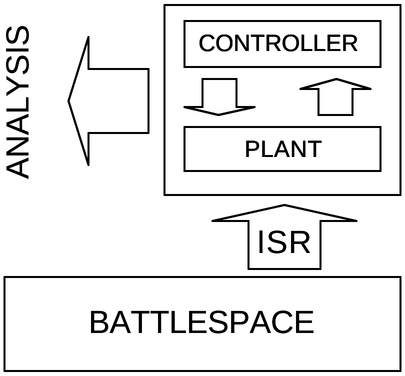

Figure 3.

Note that the controller/plant pair is now seen as a technique for performing analysis of the plant, based on the best available state estimates via ISR data. Rather than issuing control signals directly to the real plant (the battlespace), the controller operates against the on-line simulation. This approach enables a number of kinds of analysis, each capable of assisting the commander in understanding the battlespace, and focusing attention on what is most important at the moment.

For example, the commander could use this configuration to perform what-if analysis, comparing alternative strategies against a variety of metrics (timeliness, attrition, target value destroyed, etc.). This process is analogous to the algorithm employed by MPC, which internally relies on a configuration very similar to that in Figure 3.

The key here is that the controller technology is now embedded in, and drives the design of, a simulation able to drive the plant forward in time in order to assess and analyze the likely effect of decisions to be made in the immediate future.

3.3 Tracking and Adaptive Control

Another type of analysis enabled by Figure 3 is to track plant behavior against simulated behavior, and alert a commander when the residuals begin to become statistically significant. The idea here is that if the models embedded in the controller are reasonably accurate, then the behavior it predicts and the behavior the plant produces should agree to within some acceptable tolerance. As soon as plant behavior starts deviating from predicted behavior, this is a sign that the model is not longer reliable.

Now, *if* the model in fact captures the commander's current perception of the battlespace, this deviation from nominal could be used to alert the commander that his/her assumptions are no longer valid and that a shift is in order. Used in this way, the arrangement in Figure 3 in effect is a scheme to alert the commander to the possibility that things are starting to head in an unforeseen direction.

Carrying this idea one step further, this type of tracking could be the basis for on-line adaptation to changing circumstances. For example, it is becoming clear as a result of experimentation that different control technologies perform more or less well depending on specific operational characteristics. One may perform best when forces are evenly balanced, another may shine in an asymmetrical engagement. Thus, as the battle evolves and the operational picture changes, it may be best to switch from one control strategy to another.

This change would be more than just adjustment of model parameters, or switching between alternative local linearizations (as in Hybrid control). It means understanding the strengths and weaknesses of different control approaches, and selecting among them based on actual battle conditions and commander's intent. If the controller's ability to achieve its objective is thought of as its value, and the circumstances it may face is considered a set of independent variables, then imagine generating a *value landscape* for the controller – that is, an assessment of its likely performance against variations in one or another key parameters characterizing plant behavior. The ability of a configuration like Figure 3 to produce value landscapes as analytic tools has proved to be an exciting and promising development of the project.

3.4 Sensitivity Analysis

Continuing this line of thought, imagine a configuration like Figure 3 as a prosthetic device to help the commander understand which factors in the battlespace are most important at the moment. The idea here is that the battlespace is a very complex, high-dimensional system ill-suited to natural modes of human understanding.

The need, then, is to reduce this high-dimensional space to a low-dimensional space that the human being can understand and use for decision-making. This in turn means exploring this high-dimensional space, and discovering which factors are most important *now*. Off-line, in-advance analysis is almost certainly doomed to failure. The real situation the commander faces is never just like the textbook, and so hard-wired "answers" computed in advance are likely to miss just the factor that is of interest at the moment.

A configuration like that in Figure 3, with its ability to compute the value landscape (see 3.3 above), should be able to analyze which factors *in the case at hand* are most relevant, and which, on the other hand, appear to have little or no effect on the outcome. This appears to be a promising approach to on-line decision support.

4. Conclusions, and Future Directions

It should first be pointed out that this brief summary does not include all the work being done on the JFACC project. Rather, approaches have been selected based on likely interest and, unfortunately, on limitations of time and space. For example, substantial effort has been devoted to modeling the organizational impact of various C2 architectures. Another important area recognized by many researchers is the need to decompose the problem into hierarchical levels, each with its own appropriate modeling abstractions and control approach. Finally, important work has been done on state estimation that has not been discussed.

There appear to be three "lessons learned" about the problem faced here that should guide further research efforts. *First* is the importance of accurately capturing commander's intent. Simple linear objective functions based on target value will not be adequate. Non-linear coupling and the ability to express risk and temporal constraints are also required. This is even truer when an active adversary is included in the model, with the need to estimate the opponent's intent and valuations.

Second is the need for flexible, adaptive control based on an understanding of the strengths and weaknesses of the available control approaches. The ability to compute a value landscape for use in selecting among alternative approaches is key here.

Third is the need for state estimation, and particularly the ability to draw the commander's attention to those factors that are most important in the case at hand. Again, the ability to compute a value landscape, but this time in support of sensitivity analysis, is likely to be a key functionality.

Looking to the future, the military is increasingly interested in unmanned and/or autonomous entities in the battlespace. What is the right C2 approach for this new paradigm? How are human beings to synchronize with autonomous entities to achieve common objectives, and to respond to commander's intent in an efficient and coordinated manner? These appear to be the challenges on the immediate horizon, and the technologies developed under the JFACC program seek to offer a sound basis on which to build toward the future.

References:

[1] Bertsekas, D. P. and Shreve, S. E., *Stochastic Optimal Control: The Discrete-time Case*, Athena Scientific, Belmont MA, 1996

[2] Cassandros, C. G. and Lafortune, S., *Introduction to Discrete Event Systems*, Kluwer, Norwell MA, 1999

[3] Hopcroft, J. E. and Ullman, J. D., *Automata Theory, Languages, and Computation*, Addison-Wesley, Reading MA, 1979

[4] Mayne, D. Q. et al, "Constrained model predictive control: Stability and optimality," *Automatica* 36 (2000), pp 789 - 814

[5] Stillman, B, *Linguistic Geometry: From Search to Construction*, Kluwer, Norwell MA, 2000